

REMARKS

Claims 1-2, 4-8, 10-74, and 76-82 were pending. Claims 1, 27, 44, 57, 72 and 77-80 have been amended. No claims have been added or canceled. Claim 82 was erroneously numbered and has been renumbered as claim 80. Accordingly, claims 1-2, 4-8, 10-74, and 76-80 remain pending subsequent entry of the present amendment.

35 U.S.C. § 102 and § 103 Rejections

Claims 1, 10-11, 13-19, 23,25-26, 44, 48-49, 52-54, 56-57, 60, and 69-73 stand rejected under 35 U.S.C. § 102(e) as being anticipated by U.S. Patent Publication No. 2002/0027906 (hereinafter “Athreya”), or, in the alternative, these claims stand rejected under 35 U.S.C. 103(a) as being obvious over Athreya in view of U.S. Pub. No. 200410073716 (hereinafter “Boom”). Claims 2, 12, 45-46, 51, 55, 61, 67-68, 74, and 76-79 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Athreya in view of Boom. Claims 4-8, 20-22, 24, and 47 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Athreya in view of U.S. Patent 6,823,458 (hereinafter “Lee”), or, in the alternative, under 35 U.S.C. 103(a) as being obvious over Athreya in view of and further in view of Lee. Claims 27, 28, 31-43, 58, 59, and 62 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Athreya in view of U.S. Pub. No. 200410202013 (hereinafter “Dove”), or, in the alternative, under 35 U.S.C. 103(a) as being obvious over Athreya in view of Dove) and further in view of Boom. Claims 29-30 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Athreya in view of Dove and further in view of U.S. Pub. No. 2003/0065822 (hereinafter “Avery”), or, in the alternative, under 35 U.S.C. 103(a) as being obvious over Athreya in view of Dove and Boom and further in view of Avery. Claims 63-64 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Athreya in view of U.S. Patent 6,731,649 (hereinafter “Silverman”), or, in the alternative, under 35 U.S.C. 103(a) as being obvious over Athreya in view of Boom and further in view of Silverman. Claims 50, 65, and 66 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Athreya in view of U.S.

Patent 6,834,326 (hereinafter “Wang”), or, in the alternative, under 35 U.S.C. 103(a) as being obvious over Athreya in view of Boom and further in view of Wang.

Applicant appreciates the examiner’s consideration of the present matter and has carefully reviewed both the cited art and the rejections. Applicant believes the present claims recite features neither disclosed nor suggested by the cited art. Accordingly, the rejections are traversed and reconsideration is requested in view of the following comments.

The Athreya and Boom References

As noted above, claims 1, 10-11, 13-19, 23,25-26, 44, 48-49, 52-54, 56-57, 60, and 69-73 stand rejected under 35 U.S.C. § 102(e) as being anticipated by Athreya.

By way of preface, Applicant notes the presently claimed invention is directed to art which is fundamentally distinct from Athreya. In particular, the present claims are generally directed to a PCI Express based architecture. In contrast, Athreya is generally directed to virtual local area networks and bears and in no way concerns PCI Express based architectures. While Applicant believes this to be clear in the prior claims (e.g., the recited root complex when read in light of the description is believed clear), Applicant has nevertheless amended the claims to clarify the nature of the presently claimed invention. For example, claim 1 has been amended to recite the features of prior claim 77 – namely, the root complexes are PCI Express root complexes. Additionally, claim 1 has been amended to recite a shared PCI Express endpoint comprising an I/O controller. Both PCI Express root complexes and PCI Express endpoints have well understood meanings in the art which are readily distinguished from Athreya.

Turning now to the rejection of claim 1, the features of claim 1 are readily distinguished from Athreya. In both the previous Office Action of June 27, 2007, and the present Office Action, it is suggested that networks 52, 54, 56, and 74, 76, 78, and similar networks in other Figures of Athreya are equivalent to the recited root complexes.

However, claim 1 recites the root complexes are PCI Express root complexes. Applicant submits the computing networks of Athreya are clearly not equivalent to PCI Express root complexes. Neither does Athreya disclose or suggest anything concerning PCI Express root complexes. Accordingly, for at least this reason, claim 1 is not anticipated by Athreya.

On page 16 of the present Office Action, first full paragraph, claim 77 was rejected over Athreya in view of Boom. In the rejection, Boom was cited as disclosing PCI Express root complexes. It was then suggested Athreya be modified to include the Boom PCI Express root complexes within the customer networks (e.g., 74, 76, 78) of Athreya. Even were one to accept, *arguendo*, such a combination as being appropriate, claim 1 as amended recites further features which are neither disclosed nor suggested by the proposed combination. For example, claim 1 further recites the features “a shared PCI Express endpoint coupled to said shared switch, the endpoint comprising an I/O controller configured to be shared by the at least two of said root complexes.” Nowhere do any of the cited references disclose or suggest a PCI Express endpoint configured to be shared by at least two PCI Express root complexes. Such features are wholly absent from the cited art. Accordingly, claim 1 is patentably distinguishable from the combination of cited art.

In addition to the above, on page 3 of the present Office Action, it is suggested that Athreya discloses the features “wherein at least two of said root complexes do not include a dedicated network interface controller.” In the rejection it is stated:

“Athreya further teaches said root complexes do not include a dedicated network interface controller (note that the use of a network interface controller is not disclosed anywhere in the reference, i.e. each of the root complex (Customers 52,54,56) does not include an internal or integrated network interface controller within itself).”

“... note that the use of a network interface controller is not disclosed anywhere in the reference, i.e. each of the root complex (Customers 52, 54, 56) does not include an internal or integrated network interface controller within itself). Applicant concluded that network interface

controllers are ubiquitous in paragraphs 0002, 0040 and 0041. However, these paragraphs only disclose how MAC addresses being mapped and are silent regarding a network interface controller.” (Office Action, page 27)

However, Athreya discloses in FIG. 2A, FIG. 3A, and FIG. 5, that the Customer Networks have a single connection, or a single port or a single logical interface, from the Customer Networks to a common single bridge (FIG. 2A), single router (FIG. 3A), or single switch (FIG. 5). Also, in FIG. 15 of Athreya, a single connection is shown from subnet 206 associated with an engineering building to hub site 204, and likewise, a single connection is shown from subnet 208 associated with a management building to hub site 204 via the same common bus or connection. Therefore, each of the customer networks (52, 54, 56 and 74, 76, 78 and 112, 114, 116) and each of the subnets representing customer networks (206, 208) must include a network interface controller in order to prevent contention on the common single connections. It is well known in the art that an interface controller of some kind needs to be included within each of the customer networks in the disclosed systems in the Figures in order to prevent contention on a common connection or common bus.

In addition, Athreya discloses a logical interface, or the single connection discussed above, is utilized to connect customer networks to units, such as tagging unit 64 of FIG. 2A, for virtual local area network (VLAN) tagging. Examples of logical interfaces at the unit are provided in Athreya as being a T1 or Ethernet port. Therefore, an Ethernet or T1 network interface controller is needed. For example, Athreya discloses:

“The present invention relates to Virtual Local Area Networks (VLANs). VLANs allow multiple community groups to coexist within one bridge. A bridging community is identified by its VLAN ID. As described in IEEE 802.1q standard, if a system that supports VLANs receives a frame from the LAN, the frame will only be admitted to a LAN which belongs to the same community. A disadvantage of the VLAN standard 802.1q is that it does not support interactions over Wide Area Networks (WANs). It is desired to have a VLAN supported over a wide area network because it would allow the logical grouping of resources between a customer premise equipment (CPE) location

and a point of presence (POP) location.” (Athreya, paragraph 0002) (emphasis added).

“In a preferred embodiment, a unit provides VLAN ID tagging based on a logical interface at the unit. The logical interface can be a physical interface, such as a T1 or Ethernet port, multi-link combination of physical connections, such as multi-linked T1s, or virtual connections, such as a virtual circuit, or an aggregated virtual circuit. Other possible logical interfaces include DS3, CT3, and E1.” (Athreya, paragraph 0005) (emphasis added).

“FIG. 14 illustrates a system in which a unit 194 is connected to a network 196 within a node site 198 . . . In this example, the node customer is connected by the Ethernet to the unit 194 and can be serviced by VLAN bridging between the hub site and node site . . . Units 194 , 200 and 202 are transparent between the node site and the hub site for both Ethernet and in-building T1 customers. For in-building T1 customers going from node site to hub site, unit 194 tags units with the configurable VLAN ID going towards the POP backbone over the WAN link.” (Athreya, paragraph 0065) (emphasis added).

“. . . In most systems described above, VLAN tagging is based solely on the logical interface. This means the various flows going into an interface all get tagged with the same VLAN ID. It is desired to have flow-based VLAN tagging in which, in addition to the interface for the unit, flow-based tagging is done. The additional data used for the flow-based tagging include tagging based upon TCP port, tagging based upon UDP port, tagging based on destination IP address, tagging based on source IP address, and tagging based on IP subnet. FIG. 15 illustrates a situation in which traffic is sent to the unit 204 from various buildings in the company. For example, the subnet 206 is associated with an engineering building, and building subnet 208 is associated with a management building. If the management building needs faster network access than the engineering building, the flows can be tagged by the unit 204 with different VLAN IDs, and the unit 204 or other units in the network can prioritize the data from the different subnets in a different fashion. For example, the data from subnet 208 is tagged with the VLAN ID 100 and data from subnet 206 is tagged with the VLAN ID 200. If the VLAN 200 gets higher priority than the VLAN ID 100, the users of the network 206 have higher priority over the users of the network 208 . This provides a policy-based VLAN forwarding.” (Athreya, paragraph 0066) (emphasis added).

Further, Athreya teaches the use of customer premise equipment (CPE). This equipment requires connections to a network, such as the use of Ethernet or T1 lines, as disclosed by Athreya as in the following:

“... It is desired to have a VLAN supported over a wide area network because it would allow the logical grouping of resources between a customer premise equipment (CPE) location and a point of presence (POP) location.” (Athreya, paragraph 0002).

“FIG. 5 illustrates one embodiment of the present invention. In this embodiment, the customer premise equipment (CPE), includes customer networks 112, 114 and 116.” (Athreya, paragraph 0045) (emphasis added).

“FIG. 13 illustrates a scenario in which customers in the same node or building can connect to the building CPE unit 190 by either an Ethernet or by a T1. If the connection is made through a T1, then the building CPE unit 190 tags the incoming LAN traffic on a per bundle basis, and forwards the packet out to the upstream WAN bundle interface of unit 192. In the scenario of FIG. 13, VLAN bridges between the unit 190 and 192 can be done. In building T1s will connect to the unit 190 according to a number of possible protocols.” (Athreya, paragraph 0063) (emphasis added).

CPEs may be any associated equipment at a subscriber's premises, such as the buildings 206 and 208 of FIG. 15 of Athreya, and connected with a carrier's telecommunication channel(s) at a demarcation point. The demarcation point is a point established in a building or complex to separate customer premise equipment from non-customer premise equipment (e.g., telephone company equipment). These CPEs are connected to a network, such as the disclosed Ethernet and T1 connections in Athreya, with demarcation devices that are deployed at customer locations. These demarcation devices provide a clear separation between a user and the network and allow carriers to extend network visibility into the user premises. Examples of demarcation devices, which require network interface controllers, include Ethernet Demarcation Devices (EDDs), Network Termination Equipment (NTE), and Network Interface Devices (NIDs). Therefore, for all of the above reasons, Athreya teaches the customer networks

52, 54, 56, and 74, 76, 78, and similar networks in other Figures, do include a dedicated network interface controller(s). Consequently, Athreya does not disclose “a plurality of root complexes, wherein at least two of said root complexes do not include a dedicated network interface controller”.

The rejection then states that Boom teaches a root complex does not have to include a network interface controller as in the following:

“However, for clarification, the Examiner further cited reference Boom et al. (U.S. Pub. No. 200410073716) for the teaching that a root complex does not have to include an integrated network interface controller. Boom teaches a root complex needs not have a network interface controller (NIC 320) integrated within itself (Fig. 4), since NIC 320 is a stand alone network interface card which can be connected to the root complex (Fig. 4).” (Office Action, page 3).

However, Applicant notes the claim language does not recite features regarding integration of a network interface controller. In contrast, claim 1 recites the features “at least two of said root complexes do not include a dedicated network interface controller.” Further, in Boom, FIG. 4 simply discloses a prior art example of an architecture similar to that of FIG. 3 of the present application. In the Description regarding FIG. 3 of the present application it is noted:

“What should be appreciated at this point is that the server 300 as shown still requires dedicated I/O controllers 324, 326, 328 to allow each server 300 to communicate to network fabrics such as Ethernet, Fiber Channel, etc.” (Specification, paragraph 0075).

Similarly, Boom discloses a dedicated network interface controller 320 for root complex 318. Boom discloses a system for forwarding data packets from a network interface controller (NIC) 320 to one or more media decoding devices, such as video output device 322 or audio output device 325 through the switch 316. However, Boom nowhere discloses a plurality of root complexes let alone “wherein at least two of said root complexes do not include a dedicated network interface controller”. For at least

these further reasons claim 1 is patently distinct from the cited references Athreya and Boom, taken alone or in combination. Claims 27, 44, 57 and 72 are distinguishable for similar reasons.

Similar to claim 1, claims 44, 57 and 72 as amended recite the root complexes are PCI Express root complexes, and the recited endpoints are PCI Express endpoints. Given such a clarification, and the recited features wherein a PCI Express endpoint is configured to be shared by at least two of the root complexes, each of claim 44, 57, and 72 are patentably distinguishable from the cited art for at least reasons similar to those discussed above concerning claim 1.

In view of the above, each of independent claims 1, 27, 44, 57, and 72 is patentably distinguishable from the cited art, taken either singly or in combination. Accordingly, all pending claims are patentably distinct.

Sole remaining independent claim 27 stands rejected under 35 U.S.C. 103(a) as being unpatentable over Athreya in view of Dove, or, in the alternative, under over Athreya in view of Dove and further in view of Boom. It is noted that the rejection of claim 27 depends upon Athreya in a manner similar to that of the other independent claims. As Athreya includes at least the deficiencies discussed above, the rejection of claim 27 is deficient for at least the reasons discussed above. The deficiencies of Boom discussed above are also applicable to the rejection of claim 27. Similar to claim 1, claim 27 as amended recites the endpoints are PCI Express endpoints. Given such a clarification, and the recited features “wherein at least two of said plurality of operating system domains do not include a dedicated network interface controller and are configured to share said plurality of endpoints”, claim 27 is likewise believed patentably distinguishable from the combination of cited art.

In addition to the above, the dependent claim recite further features neither disclosed nor suggested by the cited art. For example, each of claim 77-80 recite the

further features “said root complex identification identifies a particular PCI Express root complex which is a source of a received packet.” Such features are nowhere disclosed or suggested in the cited art. In the present rejections, reference is made to the disclosure of Athreya wherein a VLAN ID is placed within a packet (e.g., see page 3 of the Office Action). However, a VLAN ID is clearly not equivalent to a root complex identification that identifies a particular PCI Express root complex which is a source of a received packet. A VLAN ID bears no relation to PCI Express. Rather, the VLAN ID simply serves to indicate a given packet corresponds to a particular virtual local area network and at best may be associated with a given customer network (e.g., 74). Therefore, claims 77-80 are distinguishable for at least these further reasons.

Applicant believes the application to be in condition for allowance. However, should the examiner believe issues remain, the below signed representative would appreciate, and requests, a telephone interview to facilitate a resolution.

CONCLUSION

Applicant submits the application is in condition for allowance, and an early notice to that effect is requested. If any additional fees are due, the Commissioner is authorized to charge said fees to Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C. Deposit Account No. 501505/6107-00100/RDR

Respectfully submitted,

/Rory D. Rankin/

Name Rory D. Rankin

Reg. No. 47,884

ATTORNEY FOR APPLICANT(S)